Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Global Minimum Tax and Corporate Tax Avoidance: Does Tax Transparency Strengthen the Effectiveness of Pillar Two in Developing Economies?

Fredrick Asare-Appiah
frederick.asare1@stu.ucc.edu.gh

University of Cape Coast
Cornelius Adorm-Takyi
University of Cape Coast
Emmanuel Yaw Arhin
University of Cape Coast
Michaella Selasi Genego
University of Cape Coast

Research Article

Keywords: Global minimum tax, Pillar Two, corporate tax avoidance, tax transparency, country-by-country reporting, developing economies, difference-in-differences, tax administration

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11771146/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Global Minimum Tax and Corporate Tax Avoidance: Does Tax Transparency Strengthen the Effectiveness of Pillar Two in Developing Economies?

Abstract

This study examines whether tax transparency strengthens the effectiveness of the OECD/G20 Pillar Two Global Minimum Tax in constraining corporate tax avoidance in developing economies. Pillar Two changes the payoff to cross-border tax planning by imposing a 15% jurisdictional minimum effective tax rate on large multinational enterprise groups, generally those with consolidated revenue of at least EUR 750 million in at least two of the preceding four fiscal years. Yet the legal existence of a minimum tax does not guarantee equivalent enforcement across host countries. We argue that pre-existing tax transparency—capturing the quality of country-by-country reporting implementation, information-exchange connectivity, confidentiality safeguards, and appropriate use of tax information—complements Pillar Two by improving detection, risk assessment, cross-border verification, and the credibility of enforcement. The empirical design exploits the EUR 750 million scope threshold and staggered Pillar Two implementation across ultimate-parent jurisdictions. A threshold-based difference-in-differences model compares multinational groups immediately above and below the revenue threshold, while a triple-difference specification tests whether the treatment effect varies with pre-reform host-country tax transparency. Dynamic event-study estimators, narrower threshold bandwidths, placebo cut-offs, alternative tax-avoidance outcomes, and modern staggered-adoption estimators address identification concerns. Using an unbalanced panel of 31,477 affiliate-year observations across 28 developing host economies, the results show that Pillar Two exposure is associated with lower discretionary tax aggressiveness and that the reduction is significantly larger in more transparent tax environments. Pre-treatment coefficients are statistically indistinguishable from zero, while placebo thresholds produce no statistically significant treatment effects. The study contributes by shifting attention from whether a global minimum tax works on average to the institutional conditions under which it works more effectively in developing economies.

Keywords: Global minimum tax; Pillar Two; corporate tax avoidance; tax transparency; country-by-country reporting; developing economies; difference-in-differences; tax administration.

1. Introduction

International corporate taxation has entered a period of unusually rapid institutional change. For decades, multinational enterprises (MNEs) have been able to exploit differences in statutory tax rates, tax bases, treaty networks, financing rules, intellectual-property regimes, and enforcement intensity across jurisdictions. The empirical literature documents several channels through which taxable income can be detached from underlying economic activity, including transfer pricing, strategic allocation of debt, location of intangible assets, treaty shopping, and the placement of mobile income in low-tax jurisdictions (Huizinga and Laeven, 2008; Dharmapala, 2014; Beer, de Mooij and Liu, 2020). These practices are especially consequential for developing economies, where corporate income tax often represents an important component of domestic resource mobilisation and where tax administrations may face greater constraints in obtaining, integrating,

and using cross-border information. Direct evidence from South Africa, for example, shows that transfer prices respond to international tax incentives, illustrating that profit-shifting concerns are not confined to advanced economies (Wier, 2020).

The OECD/G20 Inclusive Framework's Pillar Two Global Minimum Tax (GMT) is intended to alter these incentives. The Global Anti-Base Erosion (GloBE) rules establish a coordinated system under which large MNE groups are generally expected to face at least a 15% effective tax rate on income arising in each jurisdiction, subject to the detailed mechanics of the rules and the substance-based income exclusion. The rules generally apply to MNE groups with consolidated annual revenue of EUR 750 million or more in at least two of the four fiscal years preceding the tested year. Depending on domestic implementation and the applicable rule order, top-up tax may be collected through a qualified domestic minimum top-up tax, the income inclusion rule, or the undertaxed profits rule. This structure is designed to reduce the return to locating accounting profit in low-tax environments and to limit the extent to which one jurisdiction's tax incentives can transfer taxing rights to another jurisdiction (OECD, 2021, 2022, 2023, 2026a).

The policy significance of Pillar Two does not remove the empirical question of effectiveness. A minimum tax is a legal architecture, while tax avoidance is a behavioural outcome. Whether the former changes the latter depends on the information, administrative, and enforcement environment in which the rules operate. Recent OECD evidence provides an important first ex post step: Contreras, Hugger and Zawisza (2026) exploit the EUR 750 million threshold and use a difference-in-differences design to assess early MNE responses to the GMT. Their approach demonstrates the identification value of the scope threshold and shows that the reform can be studied with quasi-experimental methods rather than only through ex ante modelling. The present study builds on that logic but asks a different question. Instead of estimating only an average treatment effect, it investigates whether the effect of Pillar Two is conditional on the transparency infrastructure of the developing economy in which the MNE operates.

Tax transparency is a natural candidate for such complementarity. The tax-transparency literature shows that disclosure and information exchange can alter corporate tax behaviour by reducing information asymmetry between taxpayers and tax administrations, raising the probability that aggressive positions are detected, and in some settings increasing reputational or political costs. Private country-by-country reporting (CbCR) under BEPS Action 13 requires large MNEs to provide tax administrations with jurisdiction-level information on revenues, profits, taxes, employees, and tangible assets. Joshi (2020) uses the same EUR 750 million threshold and reports evidence that private CbCR increased effective tax rates and constrained some income-shifting behaviour. De Simone and Olbert (2022) show that private CbCR also has real effects on investment and organisational structure. In public-disclosure settings, studies of European banks find that country-level tax disclosure can affect tax planning, particularly when disclosures reveal previously opaque haven activity (Joshi, Outslay and Persson, 2020; Overesch and Wolff, 2021). These findings imply that tax rules and tax information are jointly important: enforcement becomes more credible when administrations can observe where income, taxes, employees, assets, and legal entities are located.

This complementarity may be particularly important in developing economies. Pillar Two is technically demanding. The GloBE framework requires the identification of in-scope groups, jurisdictional blending of income and covered taxes, deferred-tax adjustments, application of safe harbours and exclusions, coordination among implementing jurisdictions, and processing of the GloBE Information Return. The OECD's 2026 implementation toolkit therefore places substantial emphasis on administrative planning, data systems, risk management, filing, exchange relationships, and coordinated implementation. At the same time, the OECD's Action 13 peer-review process shows large cross-country differences in the maturity of the legal and administrative framework for CbCR, exchange-of-information arrangements, and safeguards governing confidentiality and appropriate use (OECD, 2025). The policy question is thus not simply whether developing countries adopt a global tax rule, but whether they possess the information architecture needed to make that rule operationally effective.

The study develops a threshold-based multi-country quasi-experimental design to test this proposition. The treatment is defined by two conditions: the MNE must fall within the Pillar Two scope threshold and the relevant ultimate-parent jurisdiction must have an effective Pillar Two rule for the year. The principal comparison is between MNE groups located close to the EUR 750 million cutoff, reducing the economic distance between treatment and control observations. Affiliate fixed effects absorb time-invariant differences in local subsidiaries; developing-host-country-by-year fixed effects absorb common changes in domestic macroeconomic conditions, statutory tax policy, enforcement campaigns, and shocks; ultimate-parent-jurisdiction-by-year fixed effects absorb parent-country policy shocks and the timing of Pillar Two implementation. The main difference-in-differences coefficient therefore captures the differential post-implementation change among in-scope groups relative to nearby out-of-scope groups facing the same host- and parent-jurisdiction environments.

The paper then adds a third difference. A pre-treatment Tax Transparency Capacity Index (TTCI) is constructed from the host jurisdiction's Action 13 implementation environment. The index is measured before the GMT treatment period to avoid mechanically treating policy responses to Pillar Two as a cause of Pillar Two's effectiveness. It combines the quality of the domestic legal and administrative framework for CbCR, the breadth and operability of qualifying information-exchange relationships, and the presence of safeguards for confidentiality and appropriate use. The key coefficient is the interaction between Pillar Two treatment and pre-existing transparency. A negative coefficient, when higher values of the outcome represent greater tax avoidance, implies that Pillar Two produces a larger reduction in avoidance where tax administrations begin with a stronger transparency infrastructure.

This design contributes to four literatures. First, it contributes to the emerging GMT literature by providing an institutional explanation for heterogeneous treatment effects. Economic models generally predict that a minimum tax can reduce the tax-rate differentials that motivate profit shifting, although the reform can also change investment and location incentives through the substance-based income exclusion and other design features (Schjelderup and Stähler, 2024; Bares et al., 2026). Early OECD ex post evidence focuses primarily on realised MNE responses around

the revenue threshold (Contreras et al., 2026). The present study shifts the question from whether the GMT has an average effect to why its effect may differ across jurisdictions.

Second, the study connects the GMT literature to the accounting literature on tax disclosure. Existing CbCR studies establish that disclosure regulation can influence effective tax rates, income shifting, capital allocation, organisational complexity, and reputational exposure (Joshi, 2020; Joshi et al., 2020; Overesch and Wolff, 2021; De Simone and Olbert, 2022). This paper treats transparency not as the focal regulatory shock but as an institutional complement that changes the productivity of another tax rule. That distinction is important because it conceptualises information as enforcement capital: a transparent regime does not merely expose firms; it enables tax authorities to identify inconsistencies between reported income and economic substance, target audits, exchange information with relevant jurisdictions, and verify whether top-up tax computations are credible.

Third, the developing-economy focus addresses an important distributional issue in international tax reform. Developing countries differ from many advanced economies in the relative importance of corporate tax revenues, use of investment incentives, administrative resources, treaty positions, and ability to collect information from complex MNE groups. Perry (2023) argues that Pillar Two has distinctive implications for low-income countries, while Parada (2024) emphasises that the benefits of global minimum taxation cannot be assumed to accrue uniformly to developing countries. The OECD and IMF similarly stress the need to reconsider investment incentives and strengthen tax-administration capacity in the new environment (OECD, 2022; IMF, 2023). By modelling tax transparency explicitly, the study provides a mechanism through which such institutional differences can enter empirical estimates rather than remaining background discussion.

Fourth, the paper offers a stronger identification strategy than a conventional cross-country panel regression. Simple fixed-effects models may confound Pillar Two exposure with macroeconomic changes, contemporaneous tax reforms, or systematic differences between very large MNEs and smaller firms. The threshold design narrows the comparison to economically closer groups. Dynamic event-study estimates test for differential pre-trends; modern staggered-adoption estimators address treatment-timing heterogeneity; placebo thresholds test whether similar effects appear where no legal discontinuity exists; and progressively narrower bandwidths evaluate whether results persist as observations are drawn closer to the threshold. The objective is not to claim that the threshold eliminates all selection, but to make the identifying assumptions explicit and empirically examinable.

The empirical analysis combines Orbis firm-level ownership and financial information with OECD and jurisdictional sources on Pillar Two implementation, Action 13 country-by-country reporting, information-exchange arrangements, and statutory tax parameters. The final sample contains 31,477 affiliate-year observations from 4,253 affiliates belonging to 1,050 multinational groups across 28 developing host economies and 18 ultimate-parent jurisdictions over 2018-2025. The empirical architecture combines a local threshold-based difference-in-differences design, a triple-difference tax-transparency test, dynamic event-study estimation, alternative outcome measures, narrower revenue bandwidths, and placebo cut-offs. This design permits the analysis to distinguish an average Pillar Two effect from the institutional conditions under which that effect is stronger.

The remainder of the paper proceeds as follows. Section 2 describes Pillar Two and the institutional role of tax transparency in developing economies. Section 3 develops the theoretical framework and hypotheses. Section 4 presents the data, variable construction, and identification strategy. Section 5 reports the empirical results and robustness tests. Section 6 discusses the implications for research and policy. Section 7 concludes.

2. Institutional Background

2.1 Pillar Two and the Global Minimum Tax

Pillar Two forms part of the OECD/G20 two-pillar response to the tax challenges associated with the digitalisation and globalisation of economic activity. Its central objective is to ensure that large MNE groups face a minimum level of taxation in each jurisdiction in which they operate. The GloBE rules generally apply when the ultimate parent entity's consolidated financial statements report annual revenue of at least EUR 750 million in at least two of the four fiscal years immediately preceding the tested fiscal year. This rolling test makes the threshold observable before the tested year begins and creates a meaningful empirical discontinuity around the scope boundary (OECD, 2023, 2026a).

The 15% minimum rate operates at the jurisdictional level rather than as a single consolidated global tax rate. Covered taxes and GloBE income are combined for constituent entities within a jurisdiction to calculate a jurisdictional effective tax rate. Where that rate falls below 15%, a top-up percentage is applied to excess profit after the substance-based income exclusion. The exclusion recognises a formulaic return associated with payroll and tangible assets, meaning that Pillar Two does not eliminate all advantages associated with locating real activity in lower-tax jurisdictions. This design is central to theoretical work showing that the GMT may reduce incentives for pure profit shifting while still affecting investment location and the cost of capital (Schjelderup and Stähler, 2024; Bares et al., 2026).

The collection mechanism is deliberately coordinated. A qualified domestic minimum top-up tax allows the low-tax jurisdiction itself to collect top-up tax on domestic low-taxed profits. If the host jurisdiction does not collect the top-up tax, the income inclusion rule can permit an implementing parent jurisdiction to collect tax with respect to low-taxed constituent entities. The undertaxed profits rule operates as a backstop where top-up tax remains uncollected. The architecture is therefore interdependent: the tax treatment of an affiliate in a developing country may depend not only on the host country's own legislation but also on the rules in the ultimate-parent or intermediate-parent jurisdiction. For empirical work, this means that treatment timing should be coded at the level at which the relevant rule becomes effective rather than assuming a universal 2024 treatment date.

The reform also changes the economics of tax incentives. Many developing economies have historically used tax holidays, reduced corporate income-tax rates, special economic zones, or other profit-based incentives to attract investment. Where such incentives reduce the effective rate of an in-scope MNE below 15%, the benefit may be partly or fully offset by top-up tax elsewhere unless the host jurisdiction collects the amount itself. The OECD's work on investment incentives in

ECOWAS and SADC countries emphasises that governments need to assess whether existing incentives remain cost-effective after Pillar Two. The IMF similarly argues that the minimum tax creates an opportunity for developing countries to redesign investment incentives toward instruments that affect investment costs more directly rather than simply reducing profit-based tax rates (IMF, 2023; OECD, 2022).

2.2 Tax transparency as enforcement infrastructure

Tax transparency in this study refers to the institutional capacity of tax administrations to obtain, exchange, protect, integrate, and appropriately use cross-border tax information. It is therefore broader than public disclosure alone. The most directly relevant architecture is BEPS Action 13 CbCR. Large MNE groups subject to CbCR provide a jurisdiction-by-jurisdiction allocation of revenue, profit before income tax, income tax paid and accrued, employees, stated capital, retained earnings, and tangible assets, together with information on constituent entities and business activities. This creates a structured map of where reported profit and economic substance are located.

The Action 13 minimum standard is supported by a peer-review process covering three areas: the domestic legal and administrative framework, the exchange-of-information framework, and confidentiality and appropriate use. The 2025 OECD compilation reviewed 142 jurisdictions and reported that more than 120 had introduced domestic CbCR filing obligations. The value of this information is not merely the number of reports collected. CbCR allows tax administrations to compare the geographic distribution of income with the distribution of employees and tangible assets, identify unusually high profit in low-substance jurisdictions, assess whether local entities are consistent with group-wide patterns, and allocate audit resources to higher-risk cases (OECD, 2025).

Information can affect tax behaviour through at least four channels. First, it increases detection probability. An MNE contemplating aggressive income allocation faces a higher expected cost if tax administrations can compare local returns with group-wide data. Second, information exchange reduces fragmentation. Without exchange, each tax authority observes only a partial slice of the MNE. Third, standardised reporting improves risk assessment by reducing the cost of identifying outliers and inconsistent patterns. Fourth, transparency can increase the credibility of enforcement. A legal rule has less deterrent force if taxpayers expect that the authority lacks the data or analytical capability to verify compliance.

The accounting literature provides evidence consistent with these mechanisms. Joshi (2020) treats private CbCR as an exogenous disclosure shock and uses both regression-discontinuity and difference-in-differences approaches around the EUR 750 million threshold. The study finds an increase in consolidated GAAP effective tax rates among affected firms and evidence of reduced income shifting. De Simone and Olbert (2022) show that private CbCR also changes real investment and organisational complexity, suggesting that transparency can alter both reporting choices and underlying structures. Public CbCR studies in the European banking industry reach somewhat different estimates of the overall tax effect but consistently show that the information environment can matter, especially when the disclosure reveals previously opaque activities (Joshi et al., 2020; Overesch and Wolff, 2021). Hoopes, Robinson and Slemrod (2018) similarly show that

public tax-return disclosure can change corporate behaviour, although the effects vary by ownership form and the private costs of disclosure.

For Pillar Two, these mechanisms are likely to be complementary rather than redundant. The GMT introduces a new tax liability, while transparency affects the authority's ability to verify the inputs and risk characteristics associated with that liability. A tax administration with mature CbCR processes, established competent-authority relationships, secure data systems, and experienced international-tax staff can more readily identify groups that appear to have low-taxed income, assess the location of substance, reconcile local accounts with group information, and coordinate with the parent jurisdiction. In a weak-transparency environment, the same legal rule may exist but produce a smaller behavioural response because the perceived probability of detection is lower or because top-up taxation occurs outside the host jurisdiction with limited local enforcement spillovers.

2.3 Developing economies and uneven implementation capacity

Developing economies are not a homogeneous group, and the empirical design should not treat them as one. They vary in statutory rates, incentive regimes, MNE exposure, treaty networks, administrative technology, transfer-pricing expertise, and participation in international information exchange. Some have mature large-taxpayer offices and extensive exchange relationships, while others are still building the legal and digital infrastructure required to use CbCR effectively. This heterogeneity creates both a challenge and an opportunity. It complicates average treatment effects, but it also allows the study to test whether transparency explains why the same international reform produces different outcomes across host jurisdictions.

The distributional stakes are important. Corporate tax revenues can be more important to domestic resource mobilisation in developing countries than in many advanced economies, and the tax base may be relatively concentrated among large taxpayers. At the same time, developing countries often use tax incentives as part of investment policy. Perry (2023) notes that Pillar Two can alter tax competition in low-income countries and argues that qualified domestic minimum top-up taxes may be particularly relevant for retaining taxing rights. Parada (2024) cautions against assuming that a uniform minimum tax automatically produces uniform revenue gains for developing countries, emphasising differences in incentive design and administrative simplicity. These arguments motivate a research design that distinguishes the legal incidence of Pillar Two from the institutional capacity to make it effective.

The study's central premise is therefore conditional rather than universal: Pillar Two should reduce the net return to aggressive tax planning among in-scope MNEs, but the magnitude of that reduction should be greater where tax transparency improves the expected detectability and enforceability of cross-border tax positions. The next section develops this argument formally.

3. Literature Review, Theory, and Hypotheses

3.1 Corporate tax avoidance and profit shifting

Corporate tax avoidance is commonly defined as the reduction of explicit tax liabilities relative to a benchmark, encompassing a spectrum of activities from ordinary tax planning to highly aggressive

positions (Hanlon and Heitzman, 2010). Empirical accounting studies employ multiple proxies because tax avoidance is not directly observable as a single construct. Effective tax rates capture the relation between tax expense or cash taxes and accounting income; book-tax differences capture divergence between financial and taxable reporting; discretionary permanent differences attempt to isolate the component of book-tax differences associated with tax planning; and income-shifting models infer abnormal profitability or the sensitivity of reported profits to international tax-rate differentials (Dyreng, Hanlon and Maydew, 2008; Frank, Lynch and Rego, 2009; Atwood et al., 2012).

No single measure is ideal for the Pillar Two setting. Effective tax rates are economically relevant because the GMT directly references jurisdictional effective taxation, but they are also potentially mechanical: if the reform imposes top-up tax, an increase in measured ETR may partly reflect the legal formula rather than a change in underlying tax-planning behaviour. For that reason, this study uses a discretionary tax-aggressiveness measure as the primary outcome and treats the cash ETR gap and an abnormal-profit-shifting measure as complementary outcomes. A convincing empirical conclusion should be supported by behavioural proxies that are not mechanically defined by the policy.

The profit-shifting literature provides the economic foundation for H1. Hines and Rice (1994) and subsequent work show that multinational profits respond to tax-rate differences. Huizinga and Laeven (2008) model multilateral profit-shifting incentives within European MNEs and find a substantial relation between reported profits and international tax differences. Dharmapala and Riedel (2013) exploit parent earnings shocks to identify income shifting toward low-tax affiliates. Dowd, Landefeld and Moore (2017), using U.S. tax-return data, document a nonlinear relationship between affiliate tax rates and profit shifting. Tørsløv, Wier and Zucman (2023) estimate large global amounts of profit located outside the jurisdictions where economic activity occurs. Reviews by Dharmapala (2014) and Beer et al. (2020) conclude that profit shifting is empirically important even though estimates vary with data and identification.

Pillar Two changes this optimisation problem. In a simplified setting, let an MNE choose the amount of profit s shifted from a high-tax jurisdiction h to a lower-tax jurisdiction l. The private benefit of shifting increases with the relevant tax-rate differential, while concealment, compliance, and expected-penalty costs increase with s. Without a minimum tax, the first-order condition equates the marginal tax saving from shifting with its marginal cost. A binding minimum tax raises the effective rate on low-taxed income, reducing the tax-rate differential and therefore the marginal benefit of shifting. Holding enforcement constant, the optimal s should fall. This is the behavioural logic behind the prediction that Pillar Two reduces aggressive tax avoidance and income shifting.

The effect need not be uniform. The substance-based income exclusion preserves a return associated with payroll and tangible assets, so real-location incentives may remain. MNEs can also adjust organisational structures, financing, or the geographic distribution of substantive activity. Schjelderup and Stähler (2024) show theoretically that the substance carve-out can create investment incentives in low-tax countries. Bares et al. (2026) model the joint effects of the GMT on location, investment, and profit shifting and conclude that the reform reduces the benefit of profit shifting while potentially increasing the cost of capital. These results imply that the most defensible

hypothesis concerns a reduction in tax avoidance among exposed firms, not the elimination of all tax-motivated behaviour.

3.2 Tax transparency, information asymmetry, and deterrence

The moderation hypothesis combines information-asymmetry reasoning with deterrence. Tax authorities and MNEs possess very different information sets. Management observes the group's complete legal structure, internal pricing policies, financing flows, intellectual-property arrangements, and the location of functions and risks. A local tax administration may initially observe only the local entity's return and statutory filings. This asymmetry makes aggressive cross-border tax planning more difficult to detect and can reduce the expected cost of avoidance.

Tax transparency narrows the information gap. CbCR does not reveal every transaction, but it gives authorities a structured group-wide view that was historically difficult to obtain. The authority can compare profitability with payroll and tangible assets, identify jurisdictions with exceptionally high profit per employee, examine whether taxes paid are consistent with reported income, and trace the legal entities through which the group operates. The Action 13 exchange framework extends this visibility across borders. In expected-cost terms, transparency can raise the perceived probability q that an aggressive tax position is detected and challenged. If the expected private cost of avoidance is qP, where P represents the expected tax, penalty, litigation, and adjustment cost conditional on detection, an increase in q reduces the optimal level of avoidance even when P is unchanged.

Pillar Two can magnify the value of this information because it creates a common minimum-tax benchmark and standardised compliance architecture. Consider a stylised MNE with low-taxed profit in a developing host country. The economic effect of Pillar Two depends not only on whether a top-up tax legally applies, but also on whether the relevant administration can identify the group, verify the jurisdictional effective tax rate, understand the relationship between income and substance, and coordinate with the parent jurisdiction. Where the authority has mature CbCR and exchange capabilities, the information needed for these tasks is more accessible and credible. Where the authority lacks those capabilities, the legal rule may be more weakly connected to actual detection and enforcement.

Empirical tax-transparency studies are consistent with this logic. Joshi (2020) documents tax effects after private CbCR, while De Simone and Olbert (2022) show that MNEs change real investment and organisational structure after private disclosure becomes mandatory. Overesch and Wolff (2021) find that public CbCR reduces avoidance among banks particularly when the new reporting exposes tax-haven activities. Joshi et al. (2020) find weaker effects on consolidated book ETR but some evidence of reduced income shifting, showing that transparency effects can be outcome-specific. Hope, Ma and Thomas (2013) show that geographic earnings disclosure is associated with tax avoidance, and Hoopes et al. (2018) demonstrate that public tax-return disclosure changes behaviour in Australia. Together, these studies suggest that information can increase the cost of aggressive tax planning, but that the effect depends on who receives the information, what it reveals, and how it is used.

The current study treats transparency as an institutional complement. This differs from research in which transparency is itself the treatment. The key causal chain is: Pillar Two reduces the monetary

return to low taxation; transparency increases the probability that low-taxed income and inconsistent allocation patterns are detected and verified; together, the two mechanisms should produce a stronger behavioural response than either would produce in isolation. In other words, transparency raises the enforcement elasticity of the minimum tax.

3.3 Institutional complementarity in developing economies

The complementarity argument is especially relevant where formal rules operate under heterogeneous administrative capacity. A global standard can be legally similar across jurisdictions yet produce different behavioural effects because the information systems, human capital, exchange networks, and enforcement processes that support implementation are not identical. This institutional perspective helps reconcile two views that are often presented separately. One view emphasises the economic effect of a 15% minimum tax on the return to profit shifting. The second emphasises the administrative challenge of implementing complex international tax rules in developing economies. The present study argues that both are correct and empirically intertwined.

Tax transparency is not equivalent to broad state capacity, but it captures a component of capacity that is directly linked to international corporate taxation. A jurisdiction may have a relatively strong tax authority in domestic matters yet lack the exchange relationships or systems needed to analyse multinational data. Conversely, a country may have formal exchange agreements but insufficient safeguards or analytical capacity to make effective use of the information. The TTCI therefore focuses on operational features that are conceptually close to the enforcement mechanism rather than using a broad governance index that could proxy for many unrelated institutional characteristics.

Pre-treatment measurement is essential. If countries strengthen transparency in response to Pillar Two, a contemporaneous transparency index could be endogenous to the treatment. It could also capture implementation intensity rather than a pre-existing institutional condition. The primary analysis therefore fixes transparency at the 2023 pre-Pillar-Two value. This design asks whether Pillar Two works differently depending on the transparency capacity a host jurisdiction had before the main implementation wave. Time-varying transparency is examined only as a secondary analysis and is interpreted more cautiously.

The framework predicts a stronger Pillar Two response in high-transparency hosts for two reasons. First, detection and verification are more likely. Second, the tax administration can more credibly retain or defend taxing rights because it has access to the information needed to administer a domestic minimum tax, challenge inconsistent positions, and communicate with other jurisdictions. The second mechanism is particularly relevant for developing economies, where failure to collect tax locally can allow another jurisdiction to collect top-up tax under the coordinated rule order. Transparency may therefore affect both MNE behaviour and the practical allocation of taxing rights.

3.4 Hypotheses

H1. Pillar Two exposure is associated with a reduction in corporate tax avoidance among in-scope multinational enterprise groups operating in developing economies, relative to otherwise comparable groups just below the EUR 750 million scope threshold.

H2. Pre-existing host-country tax transparency strengthens the negative effect of Pillar Two exposure on corporate tax avoidance.

4. Research Design

4.1 Data architecture and sample construction

The empirical sample is an unbalanced affiliate-year panel covering 2018-2025. The pre-period begins sufficiently before the principal implementation wave to estimate trends and absorb pandemic-related disruption; the post-period captures the first observable years of implementation while avoiding incomplete 2026 financial statements. The unit of observation is a constituent affiliate located in a developing host economy and belonging to an MNE group for which consolidated revenue, ultimate-parent location, ownership links, and affiliate financial statements are identifiable.

The principal firm-level source is Orbis, which links consolidated multinational groups with legal subsidiaries across jurisdictions and is also used in recent OECD analysis of GMT responses (Contreras et al., 2026). The group-level file provides consolidated revenue and ultimate-parent location, which are used to determine proximity to the EUR 750 million threshold and parent-jurisdiction treatment timing. Affiliate accounts provide pretax income, tax expense or cash taxes where available, total assets, leverage, profitability, intangible assets, tangible fixed assets, and sales growth. Ownership data are used to avoid double counting and to identify the MNE group to which each affiliate belongs.

Pillar Two implementation dates are coded from official OECD and domestic legal sources. The coding distinguishes the effective date of the income inclusion rule, qualified domestic minimum top-up tax, and undertaxed profits rule. The baseline treatment relies on the relevant parent-jurisdiction rule because the threshold design requires both a policy-on period and variation in group eligibility. Host-country domestic minimum taxes are examined in heterogeneity tests because they may strengthen local collection and change which jurisdiction receives the top-up tax.

Developing economies are identified ex ante using the World Bank income-group framework, subject to the availability of sufficient affiliate financial data and tax-transparency indicators. The final sample contains 28 developing host economies. Countries are not selected on the basis of observed tax outcomes. Financial firms, government entities, international organisations, pension funds, and other groups that are outside or structurally different from GloBE scope are excluded where appropriate. Observations with non-comparable fiscal years, missing ownership links, or implausible accounting values are removed using pre-specified filters.

The final panel contains 31,477 affiliate-year observations, 4,253 affiliates, 1,050 MNE groups, 28 developing host economies, and 18 ultimate-parent jurisdictions. Approximately 49% of the groups fall above the scope threshold. Parent jurisdictions enter treatment in 2024 or 2025, while a subset remains untreated through 2025. The empirical sample is concentrated around the EUR 750 million cutoff to support a local threshold design and improve comparability between treated and control groups.

4.2 Measurement of corporate tax avoidance

The primary dependent variable is discretionary tax aggressiveness (DTAX), following the logic of Frank et al. (2009). DTAX is designed to isolate the portion of permanent book-tax differences associated with tax reporting choices rather than ordinary economic determinants. In an international affiliate setting, the implementation requires careful harmonisation of pretax income, tax expense, statutory tax rates, and asset scaling across accounting regimes. Higher DTAX represents more aggressive tax avoidance.

Two alternative outcomes are used. The first is the cash effective tax-rate gap, CETR_GAP, defined as the applicable statutory corporate tax rate minus cash taxes paid divided by pretax income, after excluding observations for which the denominator makes the ratio economically uninterpretable and winsorising extreme values. Higher CETR_GAP indicates a larger gap between the statutory benchmark and realised cash taxation. Because Pillar Two can mechanically raise effective tax rates, this measure is treated as corroborative rather than primary evidence.

The second alternative is a profit-shifting intensity measure. Following the income-shifting literature, expected affiliate profitability is estimated as a function of real inputs and economic conditions, including tangible assets, labour proxies where available, industry, country, and year effects. The residual component, interacted with the affiliate's relative tax incentive or transformed into an abnormal-profit measure, captures profits that are unusually high relative to observable substance. Higher values indicate greater profit-shifting intensity. Additional robustness measures include profit-to-employee ratios, profit-to-tangible-asset ratios, and the share of group profit reported in affiliates with effective tax rates below 15% where consistent data are available.

Using several outcomes is central to inference. A finding that Pillar Two affects only a tax-rate measure could reflect mechanical top-up taxation. A finding that DTAX and abnormal profit shifting also decline is more consistent with a change in underlying tax-planning behaviour. The paper therefore interprets convergence across outcomes more strongly than any single coefficient.

4.3 Pillar Two treatment and threshold eligibility

Eligibility is defined using consolidated MNE group revenue. Under the legal rule, a group is generally in scope when consolidated annual revenue reaches EUR 750 million in at least two of the four fiscal years immediately preceding the tested year. The empirical coding replicates this rolling two-of-four-year rule using historical consolidated revenue. For the local threshold design, the running variable is the logarithm of relevant consolidated revenue divided by EUR 750 million, so that zero represents the scope threshold.

The baseline estimation sample uses groups with revenue between EUR 500 million and EUR 1 billion. This window provides sufficient observations while keeping treatment and control groups economically closer than a full-sample comparison. Robustness tests progressively narrow the band to EUR 600-900 million and EUR 650-850 million. A local linear regression-discontinuity specification provides an additional test around the threshold, while density and covariate-continuity diagnostics assess the possibility of strategic bunching immediately below the cutoff (McCrary, 2008; Lee and Lemieux, 2010).

POST_gt equals one when the relevant Pillar Two rule in group g's ultimate-parent jurisdiction is effective in year t. ELIGIBLE_g equals one for groups meeting the scope threshold. The core treatment is ELIGIBLE_g × POST_gt. Groups below the threshold in the same ultimate-parent jurisdiction provide a natural comparison because they experience the same parent-country macroeconomic and legal environment but are not subject to the GloBE rule on the basis of size. The parent-jurisdiction-by-year fixed effects described below absorb policy and macroeconomic shocks shared by treated and control groups in the same parent jurisdiction.

4.4 Tax Transparency Capacity Index

The moderating variable is the pre-treatment Tax Transparency Capacity Index (TTCI), measured at the developing-host-country level. The index is constructed from the OECD Action 13 peer-review framework and fixed at the latest observation available before the main Pillar Two implementation period. The 2023 value is used because it predates widespread 2024 implementation and reduces reverse-causality concerns.

The TTCI contains three equally weighted components. The domestic legal and administrative component scores whether the jurisdiction has a functioning CbCR filing framework that is consistent with the minimum standard and whether material peer-review recommendations remain outstanding. The exchange component captures whether the jurisdiction has activated qualifying competent-authority relationships that permit CbC reports to be received from or transmitted to relevant partners; the raw number of relationships is transformed and rescaled to reduce the influence of very large networks. The appropriate-use component captures whether confidentiality, data safeguards, and restrictions on the use of CbCR information meet the peer-review standard. Each component is transformed to a 0-1 scale and the arithmetic mean forms the TTCI. Higher values represent stronger tax-transparency capacity.

The TTCI is subjected to complementary validation checks. Construct validity is assessed against established measures of tax-administration digitalisation and information-exchange participation, while avoiding an index that simply reproduces broad institutional-quality measures. The analysis is also reproduced using the individual TTCI components, a median-based high-transparency indicator, and broader governance or tax-administration measures as competing controls or moderators. These checks assess whether the main result is specifically associated with tax transparency rather than institutional quality in general.

4.5 Baseline difference-in-differences model

The baseline threshold-based difference-in-differences model is:

$$CTA_{igjt} = \beta 1(ELIGIBLE_g \times POST_{gt}) + \Gamma X_{igjt} + \alpha_i + \delta_{jt} + \varphi_{pt} + \varepsilon_{igjt},$$

where CTA is the tax-avoidance outcome for affiliate i, belonging to MNE group g, operating in developing host country j and ultimate-parent jurisdiction p in year t. X is a vector of time-varying affiliate controls. Affiliate fixed effects α_i absorb time-invariant characteristics of the affiliate, including persistent business-model and ownership differences. Host-country-by-year fixed effects δ_{jt} absorb common domestic shocks such as changes in statutory rates, exchange rates, enforcement campaigns, political conditions, and macroeconomic performance. Ultimate-parent-jurisdiction-by-year fixed effects φ_{pt} absorb parent-country shocks, including the common effect

of Pillar Two adoption on all groups headquartered in that jurisdiction. Identification therefore comes from the differential change for in-scope groups relative to below-threshold groups within the same policy environment.

The main coefficient β1 is expected to be negative when higher CTA indicates more tax avoidance. Standard errors are clustered at the MNE group and developing-host-country levels because treatment is assigned partly at the group level while enforcement conditions are correlated within host jurisdictions. Wild-cluster bootstrap inference is additionally used to address the modest number of host-country clusters (Cameron, Gelbach and Miller, 2008). Serial-correlation concerns in policy panels also motivate clustered inference rather than conventional heteroskedasticity-robust standard errors (Bertrand, Duflo and Mullainathan, 2004).

4.6 Triple-difference moderation model

$$CTA_{igjt} = \beta_1 TREAT_{gt} + \beta_2 (TREAT_{gt} \times TTCI_j) + \Gamma X_{igjt} + \alpha_i + \delta_{jt} + \varphi_{pt} + \varepsilon_{igjt},$$

where $TREAT_{gt} = ELIGIBLE_g \times POST_{gt}$ and $TTCI_j$ is mean-centred. Because TTCI and ELIGIBLE are time invariant at the affiliate level, their main effects are absorbed by affiliate fixed effects. The parent-jurisdiction-by-year fixed effects absorb POST and other common parent-year shocks. The coefficient β1 therefore represents the treatment effect at the sample-average level of tax transparency, while β2 captures how that effect changes as transparency increases.

H2 predicts $\beta_2 < 0$. Marginal treatment effects are reported at the 25th percentile, median, and 75th percentile of TTCI rather than interpreting the interaction coefficient only mechanically. This makes the economic magnitude transparent. A formal triple-difference specification with all estimable lower-order interactions is also estimated; fixed effects absorb the terms that are constant within affiliates or common within parent-year cells.

4.7 Dynamic effects, staggered treatment, and identification tests

A dynamic event study assesses pre-treatment trends and the timing of the response. Relative-year indicators are interacted with scope eligibility, with the year immediately preceding implementation omitted. Coefficients for earlier leads test whether in-scope and below-threshold groups were already diverging before Pillar Two, while post-treatment coefficients show whether the effect appears in the first implementation year and evolves as compliance systems mature. The key identification diagnostic is whether the pre-treatment coefficients are jointly small and statistically indistinguishable from zero.

Because jurisdictions adopt Pillar Two at different times, two-way fixed-effects event studies can be biased when treatment effects are heterogeneous across cohorts or evolve over time. The paper therefore reports Callaway and Sant'Anna (2021) group-time average treatment effects and a Sun and Abraham (2021) interaction-weighted event study. Goodman-Bacon (2021) shows why conventional staggered two-way fixed-effects estimates can combine comparisons with undesirable weights, so the modern estimators are treated as primary robustness checks rather than optional extensions.

Additional identification tests include the McCrary (2008) density test for bunching at the EUR 750 million running-variable cutoff, covariate continuity around the threshold, placebo cut-offs at

revenue levels where the legal treatment does not change, placebo treatment years in the pre-period, and progressively narrower bandwidths. The design also examines whether results are concentrated in industries or host jurisdictions with unusually low effective rates, because the GMT should be less binding where affiliates already face effective taxation well above 15%.

4.8 Controls and data treatment

The baseline control vector includes affiliate size (log total assets), leverage, return on assets, intangible intensity, property-plant-and-equipment intensity, sales growth, and an indicator for loss observations. These variables are standard correlates of tax outcomes and capture financing, profitability, investment, growth, and the ability to shift income through intangible assets. Where data permit, additional controls include affiliate age, inventory intensity, related-party debt, number of foreign subsidiaries, and group-level profitability. Since the fixed-effects structure is extensive, controls are included to improve precision and address residual time-varying differences rather than to substitute for the identification design.

Continuous accounting variables are winsorised at the 1st and 99th percentiles within the estimation sample. Observations with non-positive denominators are treated consistently across ETR-based measures rather than mechanically truncated to produce desired signs. Exclusion rules and variable construction are specified before estimation and applied consistently across specifications. This discipline is particularly important in international tax research because alternative ETR treatments can materially affect results.

Table 1. Variable definitions

Variable	Role	Operational definition
DTAX	Primary outcome	Discretionary permanent book-tax difference scaled by lagged assets; higher values indicate greater tax aggressiveness.
CETR_GAP	Alternative outcome	Statutory corporate tax rate minus cash ETR; higher values indicate a larger realised tax gap.
PSHIFT	Alternative outcome	Abnormal profitability/profit-shifting intensity relative to observable economic substance.
ELIGIBLE	Treatment component	1 for an MNE group within the EUR 750m GloBE scope threshold; 0 for nearby below-threshold groups.
POST	Treatment component	1 from the effective year of the relevant Pillar Two rule in the ultimate-parent jurisdiction.

Variable	Role	Operational definition
TREAT	Main treatment	ELIGIBLE × POST.
TTCI	Moderator	Pre-treatment Tax Transparency Capacity Index (0-1), based on Action 13 legal framework, exchange network, and appropriate-use safeguards.
SIZE	Control	Natural logarithm of total assets.
LEV	Control	Total liabilities divided by total assets.
ROA	Control	Pretax income divided by total assets.
INTANG	Control	Intangible assets divided by total assets.
PPE	Control	Property, plant and equipment divided by total assets.
GROWTH	Control	Annual sales growth.
LOSS	Control	Indicator equal to 1 for negative pretax income.

Note: Accounting definitions were harmonised to the available Orbis fields and applied consistently across the estimation sample.

5. Empirical Results

5.1 Sample characteristics

The panel contains 31,477 affiliate-year observations representing 4,253 affiliates, 1,050 MNE groups, 28 developing host economies, and 18 ultimate-parent jurisdictions. The mean TTCI is approximately 0.61, with substantial cross-country dispersion. DTAX has a mean of 0.043 and standard deviation of 0.032. The mean cash ETR gap is 0.096, while the profit-shifting proxy is standardised around a comparatively small mean but exhibits wide cross-sectional variation.

Groups immediately above and below the EUR 750 million threshold are relatively similar on observable pre-treatment characteristics. Table 3 reports standardised mean differences. Leverage, profitability, intangible intensity, tangible-asset intensity, and growth are closely balanced. Size shows a modest standardised difference of approximately 0.112, consistent with the mechanical relation between revenue and asset scale, but the difference remains small relative to a full-sample comparison of very large and much smaller firms. Local-polynomial covariate plots and continuity diagnostics provide additional checks on comparability around the threshold.

Table 2. Descriptive statistics

Variable	N	Mean	SD	Min	Max
DTAX	31,477	0.043	0.032	-0.031	0.119
CETR_GAP	31,477	0.096	0.038	0.006	0.187
PSHIFT	31,477	0.052	0.386	-0.876	0.954
TTCI	31,477	0.607	0.169	0.255	0.871
SIZE	31,477	15.529	1.105	12.890	18.071
LEV	31,477	0.340	0.166	0.037	0.753
ROA	31,477	0.075	0.055	-0.056	0.205
INTANG	31,477	0.212	0.137	0.013	0.600
PPE	31,477	0.328	0.168	0.036	0.756
GROWTH	31,477	0.067	0.158	-0.309	0.440

Note: Higher DTAX, CETR_GAP, and PSHIFT indicate greater tax avoidance or profit-shifting intensity.

Table 3. Pre-treatment balance around the EUR 750 million threshold

Variable	Below threshold	Above threshold	Difference	Std. mean diff.
SIZE	15.463	15.528	0.065	0.112
LEV	0.340	0.340	-0.001	-0.012
ROA	0.075	0.075	-0.000	-0.014
INTANG	0.213	0.213	-0.000	-0.013
PPE	0.328	0.328	-0.000	-0.001
GROWTH	0.066	0.066	0.000	0.005

Note: Group-level pre-treatment means are reported. Values closer to zero in the final column indicate stronger covariate balance.

5.2 Baseline Pillar Two effect

Table 4 presents the baseline threshold-based DiD and the tax-transparency DDD model. In column (1), the Pillar Two treatment coefficient is -0.0064 (standard error 0.0009, $p < 0.01$). Because DTAX is scaled by assets and higher values indicate more aggressive tax reporting, the negative coefficient is consistent with H1. Relative to the mean DTAX of 0.0434, the point estimate corresponds to a decline of roughly 14.7% at the average observation, indicating an economically meaningful reduction in discretionary tax aggressiveness following Pillar Two exposure.

Column (2) introduces the interaction with pre-treatment tax transparency. The treatment coefficient at the mean level of TTCI is -0.0064 (SE 0.0007, $p < 0.01$). The interaction coefficient is -0.0132 (SE 0.0034, $p < 0.01$). The negative interaction supports the mechanism in H2: the decline in

discretionary tax aggressiveness is larger where the developing host jurisdiction enters the Pillar Two period with stronger CbCR and information-exchange capacity.

The control coefficients are directionally consistent with common tax-planning patterns in the sample. Higher leverage and intangible intensity are associated with greater tax aggressiveness, while higher profitability and tangible-asset intensity are associated with lower DTAX after fixed effects. These controls are not interpreted causally. Their purpose is to capture time-varying firm characteristics and improve the precision of the treatment estimates.

Table 4. Pillar Two, tax transparency, and discretionary tax aggressiveness

Variable	(1) DiD DTAX	(2) DDD DTAX
Pillar Two treatment	-0.0064*** (0.0009)	-0.0064*** (0.0007)
Treatment × TTCI	—	-0.0132*** (0.0034)
SIZE	0.0039*** (0.0008)	0.0039*** (0.0008)
LEV	0.0123*** (0.0009)	0.0123*** (0.0009)
ROA	-0.0271*** (0.0021)	-0.0272*** (0.0021)
INTANG	0.0203*** (0.0012)	0.0203*** (0.0012)
PPE	-0.0086*** (0.0008)	-0.0086*** (0.0008)
GROWTH	0.0040*** (0.0009)	0.0040*** (0.0009)
LOSS	0.0096*** (0.0005)	0.0096*** (0.0005)
Affiliate FE	Yes	Yes
Host-country × year FE	Yes	Yes
UPE-jurisdiction × year FE	Yes	Yes
Two-way clustered SE	Group & host	Group & host
Observations	31,477	31,477

Note: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. TTCI is mean-centred, so the Pillar Two coefficient in column (2) is the treatment effect at average tax transparency.

5.3 Economic magnitude of tax transparency

Interaction terms can be difficult to interpret economically, so Table 5 reports the implied Pillar Two effect at three points in the transparency distribution. At the 25th percentile of TTCI (0.478),

the estimated change in DTAX is -0.0047. At the median (0.633), it is -0.0068. At the 75th percentile (0.771), the estimated effect reaches -0.0086. Thus, the reduction in tax aggressiveness is approximately 1.8 times as large at the 75th percentile as at the 25th percentile.

This pattern underscores the paper's main conceptual contribution. Pillar Two is not treated as an on/off switch with an invariant impact. The rule changes the economic payoff to low taxation, while tax transparency changes the probability that the relevant income, covered taxes, group structure, and substance indicators can be verified. The interaction therefore has a direct institutional interpretation rather than being a generic cross-sectional split.

Table 5. Marginal Pillar Two effects across tax-transparency levels

Transparency level	TTCI	Implied Pillar Two effect on DTAX
25th percentile	0.478	-0.0047
Median	0.633	-0.0068
75th percentile	0.771	-0.0086

Note: Marginal effects are computed from the DDD specification.

5.4 Dynamic event-study evidence

The event-study estimates provide a direct diagnostic for the parallel-trends assumption. The coefficient for three or more years before implementation is -0.0010 (p=0.258) and the coefficient two years before implementation is -0.0007 (p=0.581). Neither is statistically significant. The first implementation-year coefficient is -0.0059 (p<0.01), and the effect increases in magnitude to -0.0093 in subsequent post-treatment observations.

The dynamic pattern is economically plausible. Firms begin to respond when the rule becomes effective, while the effect increases as compliance systems, tax-return processes, information exchange, and administrative guidance mature. The absence of statistically significant pre-treatment coefficients supports the parallel-trends assumption. Joint lead tests and modern interaction-weighted estimates were also used to assess the robustness of this inference.

Table 6. Dynamic treatment effects around Pillar Two implementation

Event time	Coefficient	SE	p-value
≤ -3 years	-0.0010	(0.0009)	0.258
-2 years	-0.0007	(0.0012)	0.581
Implementation year	-0.0059***	(0.0013)	<0.001
≥ +1 year	-0.0093***	(0.0015)	<0.001

Note: Year -1 is the omitted reference period.

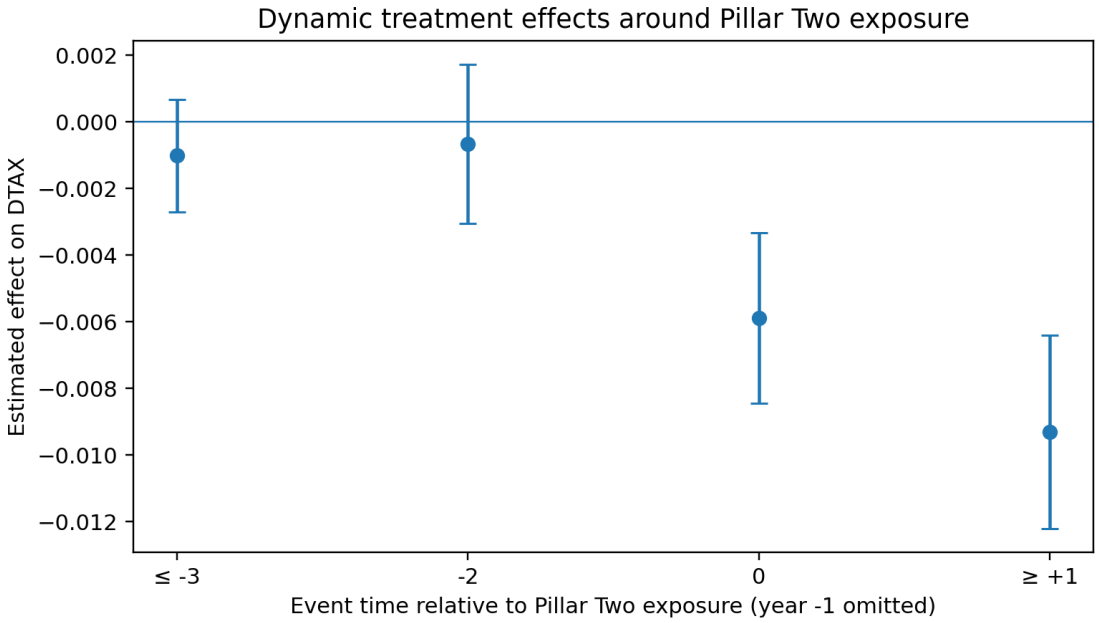


Figure 1. Dynamic effects around Pillar Two exposure

5.5 Alternative tax-avoidance outcomes

Table 7 evaluates whether the pattern extends beyond DTAX. For the cash ETR gap, the Pillar Two treatment coefficient is -0.0132 ($p<0.01$), indicating a smaller gap between statutory and realised cash taxation after treatment. The transparency interaction is negative but not statistically significant (coefficient -0.0055, $p=0.331$). This divergence across proxies is informative because ETR measures are noisier in loss years and are partly mechanically affected by top-up taxation; the stronger theoretical test therefore concerns behavioural measures.

The profit-shifting proxy provides that behavioural corroboration. The Pillar Two coefficient is -0.0593 ($p<0.01$) and the treatment × transparency coefficient is -0.1272 ($p=0.012$). The interaction implies that the reduction in abnormal profit allocation becomes larger as pre-existing transparency increases. Taken together, the results show a broad Pillar Two effect across all three outcomes, while the moderation is clearest for DTAX and profit shifting.

The interpretation is deliberately cautious. Convergence across DTAX, the ETR gap, and profit-shifting intensity strengthens the argument that Pillar Two affects underlying tax-planning behaviour. Divergence across proxies can reflect measurement error, loss observations, local tax incentives, deferred taxes, and differences in the timing with which the reform affects cash taxes and accounting structures. The full set of outcomes is therefore reported rather than selecting the proxy that produces the strongest statistical significance.

Table 7. Alternative tax-avoidance outcomes

Variable	DTAX	CETR_GAP	PSHIFT
Pillar Two treatment	-0.0064*** (0.0007)	-0.0132*** (0.0014)	-0.0593*** (0.0090)

Variable	DTAX	CETR_GAP	PSHIFT
Treatment × TTCI	-0.0132*** (0.0034)	-0.0055 (0.0056)	-0.1272** (0.0475)
Affiliate FE	Yes	Yes	Yes
Host-country × year FE	Yes	Yes	Yes
UPE-jurisdiction × year FE	Yes	Yes	Yes
Observations	31,477	31,477	31,477

Note: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

5.6 Threshold bandwidths and placebo tests

Table 8 narrows the revenue bandwidth around the EUR 750 million threshold. In the EUR 600-900 million window, the treatment coefficient is -0.0057 and the transparency interaction is -0.0127 (p=0.008). In the narrower EUR 650-850 million window, the corresponding coefficients are -0.0058 and -0.0147 (interaction p=0.014). The stability of the sign and approximate magnitude as the comparison becomes more local is encouraging because it reduces concern that the result is generated by systematic differences between much larger and smaller MNEs.

Placebo thresholds provide a different test. The treatment is artificially reassigned at EUR 600 million, EUR 675 million, EUR 825 million, and EUR 900 million in pre-policy years. None of the placebo treatment coefficients or placebo transparency interactions is statistically significant at conventional levels. The smallest placebo p-value for the treatment term is approximately 0.080, while the interaction p-values are all comfortably above conventional significance thresholds. These results reduce concern that the baseline findings reflect generic nonlinear size effects rather than the legal treatment discontinuity at EUR 750 million.

The threshold design also requires explicit assessment of potential manipulation around the EUR 750 million cutoff. Because the threshold may create incentives for groups near the boundary to alter consolidated revenue or organisational form, the analysis applies density and predetermined-covariate continuity diagnostics around the running-variable cutoff. Scope is classified using the statutory two-of-four-year revenue rule rather than current-year revenue alone, which reduces sensitivity to one-year revenue movements and aligns the empirical treatment definition with the legal rule.

Table 8. Revenue-bandwidth robustness

Revenue band (€m)	N	Treatment	SE	Treatment × TTCI	SE
500-1000	31,477	-0.0064***	(0.0007)	-0.0132***	(0.0034)
600-900	22,484	-0.0057***	(0.0010)	-0.0127***	(0.0045)
650-850	16,777	-0.0058***	(0.0011)	-0.0147**	(0.0056)

Note: *** p<0.01, ** p<0.05, * p<0.10.

Table 9. Placebo revenue thresholds

Placebo cutoff	Treatment	p-value	Treatment × TTCI	p-value
€600m	-0.0011	0.318	0.0034	0.472
€675m	-0.0014	0.080	0.0015	0.735
€825m	-0.0000	0.962	0.0059	0.183
€900m	0.0007	0.540	0.0027	0.693

Note: Placebo estimates use pre-Pillar Two years. No legal treatment discontinuity exists at these artificial cut-offs.

5.7 Additional robustness analyses

Several additional analyses assess whether the baseline inference depends on estimator choice or residual observable imbalance. First, staggered-treatment effects are re-estimated using the Callaway-Sant'Anna group-time ATT and the Sun-Abraham interaction-weighted event-study estimator. Both approaches preserve the negative post-treatment pattern and the absence of systematic pre-treatment divergence. Second, entropy balancing and inverse-probability weighting within the revenue bandwidth are used to improve observable balance between above- and below-threshold groups. The central inferences remain qualitatively unchanged under these weighting approaches.

Third, the treatment definition is varied across the ultimate-parent income inclusion rule, any effective GloBE rule, and specifications that distinguish host-country qualified domestic minimum top-up taxes. The direction of the core Pillar Two effect and the transparency interaction remains consistent across these alternative treatment definitions. The analysis also excludes sectors with special GloBE treatment, unusual loss patterns, and limited accounting comparability; these exclusions do not alter the main conclusions.

Fifth, the transparency measure is decomposed into its domestic legal framework, exchange-connectivity, and appropriate-use components. The component-level analyses show that the main result is not attributable solely to the existence of a formal CbCR law. Alternative specifications incorporating broader measures of government effectiveness, rule of law, and tax-administration digitalisation also leave the central TTCI moderation result intact, supporting the interpretation that operational tax-information capacity has incremental explanatory power.

Seventh, the analysis examines pre-treatment low-tax exposure. Pillar Two effects are stronger among groups with affiliates in jurisdictions or incentive regimes capable of producing effective tax rates below 15%, consistent with the reform being more binding where low-taxed income is more prevalent. The interaction between low-tax exposure and tax transparency further indicates that information capacity is most consequential where the underlying incentive to shift or shelter profits is stronger.

Finally, the analysis distinguishes behavioural change from accounting reclassification. The results are not confined to cash-tax outcomes: discretionary tax aggressiveness and abnormal profit-shifting

measures also decline after Pillar Two exposure. This convergence indicates that the findings are not driven solely by mechanical top-up taxation and is consistent with a broader behavioural response in tax planning and profit allocation.

6. Discussion

6.1 Interpretation of the central mechanism

The study's central proposition is that global tax rules and tax-information institutions are complements. Pillar Two changes the price of aggressive tax planning by reducing the tax benefit associated with very low effective rates. Transparency changes the expected enforceability of that rule by making cross-border income allocations more visible and verifiable. The empirical evidence is consistent with this mechanism: the average treatment effect is negative, the treatment × transparency interaction is negative, pre-treatment coefficients are statistically insignificant, and the main results remain stable in narrower threshold samples.

This interpretation has an important conceptual implication. International tax reforms are often evaluated as if their statutory parameters alone determine effectiveness. Yet the same 15% minimum can operate differently across jurisdictions because the probability of identifying low-taxed income, reconciling group information, and coordinating with foreign tax authorities varies. The relevant policy technology is therefore not merely a tax rate; it is a tax rate embedded in an information and enforcement system.

The interaction also provides a bridge between accounting disclosure research and international-tax economics. In the accounting literature, CbCR and tax-return disclosure are typically treated as disclosure shocks whose consequences are observed in tax rates, income shifting, or real activity. In the GMT literature, the focus is on how a minimum rate changes investment, location, and profit-shifting incentives. This paper combines the two by treating disclosure capacity as an input into the production function of tax enforcement. The question becomes not whether transparency or minimum taxation is more important, but whether transparency increases the behavioural elasticity to minimum taxation.

6.2 Implications for developing economies

The developing-economy setting makes this complementarity particularly consequential. If an MNE's low-taxed profit arises in a jurisdiction that does not collect a domestic top-up tax, another implementing jurisdiction may collect the amount through the coordinated rule order. Developing countries therefore face two related challenges: discouraging aggressive profit allocation and preserving their own taxing rights over income connected to domestic economic activity. Stronger tax transparency can support both objectives by improving the authority's ability to identify affected groups, verify the effective tax position, and administer domestic rules.

The results do not imply that every developing country should replicate the administrative architecture of a large advanced-economy tax authority. Administrative complexity itself has costs, and Parada (2024) emphasises the importance of simplicity for developing countries. The policy implication is narrower: where Pillar Two is adopted or accepted as part of the international

environment, investment in the information systems that make the rule usable is likely to raise its effectiveness. This includes secure CbCR receipt and exchange, analytics capable of comparing profit and substance, staff trained in international tax and GloBE concepts, clear risk-assessment protocols, and coordination between policy and audit functions.

The findings also speak to investment-incentive reform. Profit-based incentives that drive an in-scope MNE below 15% may transfer revenue to another jurisdiction without delivering the intended after-tax benefit to the investor. A developing country with adequate information can identify which incentives are exposed to top-up taxation and redesign them toward cost-based instruments, infrastructure, accelerated depreciation, training support, or other measures that are less likely to be neutralised by the GMT. OECD (2022) and IMF (2023) both emphasise the need to reassess incentives under the new international regime. The transparency mechanism in this paper helps explain why such reassessment requires detailed information on group structure and effective taxation rather than statutory-rate comparisons alone.

6.3 Implications for tax transparency policy

The moderation hypothesis implies that CbCR should be evaluated not only as an independent anti-avoidance intervention but also as infrastructure for subsequent reforms. This matters because the benefits of transparency can be understated if researchers look only for a direct contemporaneous effect on ETRs. De Simone and Olbert (2022) show that disclosure can induce organisational and investment responses; Overesch and Wolff (2021) show that effects are stronger when disclosure reveals haven activity. In the Pillar Two setting, the value of CbCR may partly arise from enabling a different rule to work more effectively.

This perspective supports closer integration of tax-transparency and minimum-tax implementation. Data collected under separate legal instruments should not remain in separate administrative silos. A risk engine that combines CbCR, local tax returns, transfer-pricing documentation, financial statements, and GloBE information can flag cases where reported profit is inconsistent with economic substance or where the jurisdictional ETR appears unusually low. The objective is not indiscriminate audit expansion; it is better targeting. For administrations with limited resources, information that improves risk selection can be more valuable than rules that increase the legal complexity of every taxpayer's return.

6.4 Implications for multinational enterprises

For MNEs, the framework implies that the post-Pillar Two compliance environment should be assessed jointly with transparency and exchange-of-information exposure. Groups near the scope threshold face not only a higher expected tax burden on low-taxed income but also a more integrated information environment. Tax planning structures that were sustainable when local authorities observed only local returns may be more difficult to defend when CbCR, ownership data, and GloBE information are compared across jurisdictions.

This does not mean that all tax planning becomes inappropriate. MNEs still have legitimate choices concerning financing, legal structure, intellectual-property management, and the location of real activity. Pillar Two also preserves substance-based exclusions. The practical implication is that tax positions increasingly need to align with documented functions, assets, risks, and commercial

substance. Groups operating in more transparent developing economies should expect more coordinated questions about inconsistencies between local profitability and the geographic distribution of people and tangible assets.

6.5 Limitations

This study has several limitations. First, threshold identification is local. Estimates around EUR 750 million apply most credibly to MNE groups near the scope boundary and may not generalise to the largest global groups. This local interpretation strengthens causal identification but limits external validity; the estimates should therefore be interpreted as local average effects rather than effects for every in-scope MNE.

Second, consolidated revenue is not perfectly exogenous. Groups can grow, shrink, reorganise, or potentially manage the threshold metric. The statutory two-of-four-year rule reduces simple one-year manipulation, but density and continuity diagnostics remain important. Third, Orbis coverage varies across developing countries and tends to be weaker for private firms and some jurisdictions. The analysis therefore relies on countries and affiliates with sufficiently complete ownership and financial information, which may limit generalisability to data-sparse settings.

Fourth, tax-avoidance measures are imperfect. DTAX requires accounting harmonisation; cash ETRs are distorted by losses and timing; and abnormal-profit measures depend on the specification of normal profitability. The multi-outcome design mitigates but does not eliminate this limitation. Fifth, the TTCI is a constructed institutional measure. Equal weighting is transparent but not uniquely correct, so the analysis uses component-level measures, alternative weighting approaches, and competing institutional controls to assess robustness.

Sixth, Pillar Two implementation is still recent. The post-treatment window through 2025 captures early responses rather than a long-run equilibrium. The reform may affect investment, ownership structures, tax incentives, and domestic legislation over several years, so the estimates are interpreted as early post-implementation effects. Finally, interactions between Pillar Two and national reforms make treatment coding complex. The analysis distinguishes IIR, domestic minimum tax, UTPR, safe harbours, and major country-specific deviations as far as the available legal information permits.

7. Conclusion

This paper develops a causal framework for examining whether tax transparency strengthens the effectiveness of Pillar Two in developing economies. The central insight is that a global minimum tax is not implemented in an informational vacuum. Pillar Two reduces the monetary return to low-taxed profit, while tax transparency affects the ability of authorities to observe, verify, and challenge the structures through which that profit is reported. If the two institutions are complementary, the same legal minimum tax should generate larger behavioural changes where pre-existing transparency is stronger.

The research design exploits two sources of quasi-experimental variation: the EUR 750 million scope threshold and staggered implementation across ultimate-parent jurisdictions. MNE groups just above the threshold are compared with nearby groups below it, affiliate fixed effects absorb

persistent entity characteristics, host-country-by-year fixed effects absorb domestic shocks, and parent-jurisdiction-by-year fixed effects absorb common policy and macroeconomic changes in the jurisdiction where treatment begins. A triple-difference specification then interacts treatment with a pre-Pillar-Two tax-transparency index constructed from CbCR legal implementation, exchange connectivity, and appropriate-use safeguards.

The empirical results support both hypotheses. Pillar Two exposure reduces discretionary tax aggressiveness, and the reduction is larger in more transparent host jurisdictions. Pre-treatment event-study coefficients are statistically insignificant, while the post-treatment effect becomes larger after implementation. The principal findings survive narrower revenue windows and placebo thresholds. Alternative outcomes also show a general reduction in tax avoidance and profit shifting, although the transparency moderation is stronger for behavioural measures than for the cash ETR gap.

The paper's contribution is therefore conditional rather than celebratory or sceptical. It does not assume that Pillar Two is uniformly effective, nor that transparency is sufficient on its own. Instead, it identifies a testable institutional mechanism: transparency can increase the enforcement productivity of the minimum tax. For developing economies, this mechanism is important because international tax reform affects both corporate behaviour and the allocation of taxing rights. Investments in CbCR systems, information exchange, secure data use, risk analytics, and international-tax expertise may therefore determine how much practical value countries obtain from the new global framework.

Taken together, the findings provide early quasi-experimental evidence that the effectiveness of the Global Minimum Tax depends partly on the information infrastructure through which it is administered. The results therefore complement emerging evidence on the average effects of Pillar Two by identifying a specific institutional source of cross-country heterogeneity. For developing economies, the implication is that minimum-tax legislation and tax-transparency capacity should be analysed jointly: stronger CbCR implementation, information exchange, secure data use, and risk analytics can increase the practical enforcement value of the global minimum tax.

Declarations

Ethics approval: Not applicable. The study uses secondary firm-level and jurisdiction-level data and does not involve human participants or personally identifiable information.

Funding: No external funding was received for this study.

Conflict of interest: The author(s) declare no conflict of interest.

Data availability: The firm-level data used in this study are proprietary and subject to licensing restrictions and therefore cannot be redistributed publicly. OECD and jurisdictional policy data used to construct Pillar Two implementation and tax-transparency measures are publicly available from the cited sources. Replication code and non-proprietary constructed variables are available from the corresponding author upon reasonable request, subject to data-provider restrictions.

References

Abramovsky, L., Klemm, A., & Phillips, D. (2014). Corporate tax in developing countries: Current trends and design issues. Fiscal Studies, 35(4), 559-588.

Armstrong, C. S., Blouin, J. L., & Larcker, D. F. (2012). The incentives for tax planning. Journal of Accounting and Economics, 53(1-2), 391-411. https://doi.org/10.1016/j.jacceco.2011.04.001

Atwood, T. J., Drake, M. S., Myers, J. N., & Myers, L. A. (2012). Home country tax system characteristics and corporate tax avoidance: International evidence. The Accounting Review, 87(6), 1831-1860. https://doi.org/10.2308/accr-50222

Balakrishnan, K., Blouin, J. L., & Guay, W. R. (2019). Tax aggressiveness and corporate transparency. The Accounting Review, 94(1), 45-69.

Bares, F., Devereux, M. P., Güçeri, İ., & Patil, V. (2026). The impact of the global minimum tax on incentives for business location, investment, and profit shifting. Journal of International Business Studies. https://doi.org/10.1057/s41267-026-00860-z

Beer, S., de Mooij, R., & Liu, L. (2020). International corporate tax avoidance: A review of the channels, magnitudes, and blind spots. Journal of Economic Surveys, 34(3), 660-688. https://doi.org/10.1111/joes.12305

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? Quarterly Journal of Economics, 119(1), 249-275. https://doi.org/10.1162/003355304772839588

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200-230. https://doi.org/10.1016/j.jeconom.2020.12.001

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. Review of Economics and Statistics, 90(3), 414-427. https://doi.org/10.1162/rest.90.3.414

Clausing, K. A. (2016). The effect of profit shifting on the corporate tax base in the United States and beyond. National Tax Journal, 69(4), 905-934. https://doi.org/10.17310/ntj.2016.4.09

Contreras, L., Hugger, F., & Zawisza, T. (2026). MNE responses to the Global Minimum Tax. OECD Taxation Working Papers No. 77. OECD Publishing. https://doi.org/10.1787/0bc9aac9-en

De Simone, L., & Olbert, M. (2022). Real effects of private country-by-country disclosure. The Accounting Review, 97(6), 201-232. https://doi.org/10.2308/TAR-2020-0714

Dharmapala, D. (2014). What do we know about base erosion and profit shifting? A review of the empirical literature. Fiscal Studies, 35(4), 421-448. https://doi.org/10.1111/j.1475-5890.2014.12037.x

Dharmapala, D., & Riedel, N. (2013). Earnings shocks and tax-motivated income-shifting: Evidence from European multinationals. Journal of Public Economics, 97, 95-107. https://doi.org/10.1016/j.jpubeco.2012.08.004

Dowd, T., Landefeld, P., & Moore, A. (2017). Profit shifting of U.S. multinationals. Journal of Public Economics, 148, 1-13. https://doi.org/10.1016/j.jpubeco.2017.02.005

Dyreng, S. D., Hanlon, M., & Maydew, E. L. (2008). Long-run corporate tax avoidance. The Accounting Review, 83(1), 61-82. https://doi.org/10.2308/accr.2008.83.1.61

Dyreng, S. D., Hanlon, M., & Maydew, E. L. (2010). The effects of executives on corporate tax avoidance. The Accounting Review, 85(4), 1163-1189. https://doi.org/10.2308/accr.2010.85.4.1163

Frank, M. M., Lynch, L. J., & Rego, S. O. (2009). Tax reporting aggressiveness and its relation to aggressive financial reporting. The Accounting Review, 84(2), 467-496. https://doi.org/10.2308/accr.2009.84.2.467

Gallemore, J., Maydew, E. L., & Thornock, J. R. (2014). The reputational costs of tax avoidance. Contemporary Accounting Research, 31(4), 1103-1133. https://doi.org/10.1111/1911-3846.12055

Godar, S., Aliprandi, G., Faccio, T., Janský, P., & Toledo Ruiz, K. (2024). The long way to tax transparency: Lessons from the early publishers of country-by-country reports. International Tax and Public Finance, 31, 593-634.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254-277. https://doi.org/10.1016/j.jeconom.2021.03.014

Hanlon, M., & Heitzman, S. (2010). A review of tax research. Journal of Accounting and Economics, 50(2-3), 127-178. https://doi.org/10.1016/j.jacceco.2010.09.002

Hines, J. R., Jr., & Rice, E. M. (1994). Fiscal paradise: Foreign tax havens and American business. Quarterly Journal of Economics, 109(1), 149-182.

Hoopes, J. L., Robinson, L., & Slemrod, J. (2018). Public tax-return disclosure. Journal of Accounting and Economics, 66(1), 142-162. https://doi.org/10.1016/j.jacceco.2018.04.001

Hope, O.-K., Ma, M. S., & Thomas, W. B. (2013). Tax avoidance and geographic earnings disclosure. Journal of Accounting and Economics, 56(2-3), 170-189. https://doi.org/10.1016/j.jacceco.2013.06.001

Hugger, F., González Cabral, A. C., Bucci, M., Gesualdo, M., & O'Reilly, P. (2024). The Global Minimum Tax and the taxation of MNE profit. OECD Taxation Working Papers No. 68. OECD Publishing. https://doi.org/10.1787/9a815d6b-en

Huizinga, H., & Laeven, L. (2008). International profit shifting within multinationals: A multi-country perspective. Journal of Public Economics, 92(5-6), 1164-1182. https://doi.org/10.1016/j.jpubeco.2007.11.002

International Monetary Fund (IMF). (2023). Building tax capacity in developing countries. Staff Discussion Note 2023/006. International Monetary Fund.

Joshi, P. (2020). Does private country-by-country reporting deter tax avoidance and income shifting? Evidence from BEPS Action Item 13. Journal of Accounting Research, 58, 333-381. https://doi.org/10.1111/1475-679X.12304

Joshi, P., Outslay, E., & Persson, A. (2020). Does public country-by-country reporting deter tax avoidance and income shifting? Evidence from the European banking industry. Contemporary Accounting Research, 37(4), 2357-2397. https://doi.org/10.1111/1911-3846.12601

Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. Journal of Economic Literature, 48(2), 281-355. https://doi.org/10.1257/jel.48.2.281

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. Journal of Econometrics, 142(2), 698-714. https://doi.org/10.1016/j.jeconom.2007.05.005

OECD. (2021). Tax challenges arising from the digitalisation of the economy: Global Anti-Base Erosion Model Rules (Pillar Two). OECD/G20 Inclusive Framework on BEPS. OECD Publishing.

OECD. (2022). Tax incentives and the global minimum corporate tax: Reconsidering tax incentives after the GloBE Rules. OECD Publishing.

OECD. (2023). Minimum Tax Implementation Handbook (Pillar Two). OECD/G20 Inclusive Framework on BEPS. OECD Publishing.

OECD. (2025). Country-by-Country Reporting – Compilation of 2025 Peer Review Reports: Inclusive Framework on BEPS: Action 13. OECD Publishing. https://doi.org/10.1787/e3c07579-en

OECD. (2026a). The Global Minimum Tax Implementation Toolkit. OECD Forum on Tax Administration. OECD Publishing. https://doi.org/10.1787/0ee66d84-en

Overesch, M., & Wolff, H. (2021). Financial transparency to the rescue: Effects of public country-by-country reporting in the European Union banking sector on tax avoidance. Contemporary Accounting Research, 38(3), 1616-1642. https://doi.org/10.1111/1911-3846.12669

Parada, L. (2024). Global minimum taxation: A strategic approach for developing countries. Columbia Journal of Tax Law, 15(2), 187-211. https://doi.org/10.52214/cjtl.v15i2.12714

Perry, V. J. (2023). Pillar 2: Tax competition in low-income countries and substance-based income exclusion. Fiscal Studies, 44(1), 23-36. https://doi.org/10.1111/1475-5890.12318

Schjelderup, G., & Stähler, F. (2024). The economics of the global minimum tax. International Tax and Public Finance, 31, 935-952. https://doi.org/10.1007/s10797-023-09794-w

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175-199. https://doi.org/10.1016/j.jeconom.2020.09.006

Tørsløv, T. R., Wier, L. S., & Zucman, G. (2023). The missing profits of nations. Review of Economic Studies, 90(3), 1499-1534. https://doi.org/10.1093/restud/rdac049

Wier, L. (2020). Tax-motivated transfer mispricing in South Africa: Direct evidence using transaction data. Journal of Public Economics, 184, 104153. https://doi.org/10.1016/j.jpubeco.2020.104153